

The Asymptotic Order of the Random k -SAT Threshold

Dimitris Achlioptas

Microsoft Research, Redmond, Washington
 optas@microsoft.com

Cristopher Moore*

Computer Science Department, University of New Mexico, Albuquerque
 and the Santa Fe Institute, Santa Fe, New Mexico
 moore@cs.unm.edu

Abstract

Form a random k -SAT formula on n variables by selecting uniformly and independently $m = rn$ clauses out of all $2^k \binom{n}{k}$ possible k -clauses. The Satisfiability Threshold Conjecture asserts that for each k there exists a constant r_k such that, as n tends to infinity, the probability that the formula is satisfiable tends to 1 if $r < r_k$ and to 0 if $r > r_k$. It has long been known that $2^k/k < r_k < 2^k$. We prove that $r_k > 2^{k-1} \ln 2 - d_k$, where $d_k \rightarrow (1 + \ln 2)/2$. Our proof also allows a blurry glimpse of the “geometry” of the set of satisfying truth assignments.

1. Introduction

Satisfiability has received a great deal of study as the canonical NP-complete problem. In the last twenty years some of this work has been devoted to the study of randomly generated formulas and the performance of satisfiability algorithms on them. Among the many proposed distributions for generating satisfiability instances, random k -SAT has received the lion’s share of attention.

For some canonical set V of n Boolean variables, let $C_k = C_k(V)$ denote the set of all $2^k \binom{n}{k}$ possible disjunctions of k distinct, non-complementary literals from V (k -clauses). A random k -SAT formula $F_k(n, m)$ is formed by selecting uniformly, independently, and with replacement m clauses from C_k and taking their conjunction¹. We will be interested in random formulas as n grows. In particular, we will say that a sequence of random events \mathcal{E}_n occurs with high probability (w.h.p.) if $\lim_{n \rightarrow \infty} \Pr[\mathcal{E}_n] = 1$.

There are at least two reasons for the popularity of random k -SAT. The first reason is that while random k -SAT in-

stances are trivial to generate they appear very hard to solve, at least for some values of the distribution parameters. The second reason is that the underlying formulas appear to enjoy a number of intriguing mathematical properties, including 0-1 laws and a form of expansion.

The mathematical investigation of random k -SAT began with the work of Franco and Paull [13]. Among other results, they observed that $F_k(n, m = rn)$ is w.h.p. unsatisfiable if $r > 2^k \ln 2$. To see this, fix any truth assignment and observe that a random k -clause is satisfied by it with probability $1 - 2^{-k}$. Therefore, the expected number of satisfying truth assignments of $F_k(n, m = rn)$ is $[2(1 - 2^{-k})^r]^n = o(1)$ for $r > 2^k \ln 2$. Shortly afterwards, Chao and Franco [6] complemented this result by proving that for all $k \geq 3$, if $r < 2^k/k$ then the following linear-time algorithm, called UNIT CLAUSE (UC), finds a satisfying truth assignment with probability at least $\varepsilon = \varepsilon(r) > 0$:

If there exist unit clauses, pick one randomly and satisfy it;
 else pick a random unset variable and set it to 0.

A seminal result in the area was established a few years later by Chvátal and Szemerédi [8]. Extending the work of Haken [18] and Urquhardt [27] they proved the following: for all $k \geq 3$, if $r > 2^k \ln 2$, then w.h.p. $F_k(n, rn)$ is unsatisfiable and every resolution proof of its unsatisfiability must contain at least $(1 + \epsilon)^n$ clauses, for some $\epsilon = \epsilon(k, r) > 0$.

Random k -SAT owes a lot of its popularity to the experimental work of Selman, Mitchell and Levesque [24] who considered the performance of a number of practical algorithms on random 3-SAT instances. Across different algorithms, their experiments consistently drew the following

*Supported by NSF grant PHY-0071139, the Sandia University Research Program, and Los Alamos National Laboratory.

¹In fact, our discussion and results hold in all common models for random k -SAT, e.g. when clause replacement is not allowed and/or when each k -clause is formed by selecting k literals uniformly at random with replacement.

picture: for $r < 4$, a satisfying truth assignment can be found easily for almost all formulas; for $r > 4.5$, almost all formulas are unsatisfiable; and for $r \approx 4.2$, a satisfying truth assignment can be found for roughly half the formulas, while the observed computational effort is maximized. The following conjecture, formulated independently by a number of researchers, captures the suggested 0-1 law:

Satisfiability Threshold Conjecture *For each $k \geq 2$, there exists a constant r_k such that*

$$\lim_{n \rightarrow \infty} \Pr[F_k(n, rn) \text{ is satisfiable}] = \begin{cases} 1 & \text{if } r < r_k \\ 0 & \text{if } r > r_k \end{cases}.$$

The conjecture was settled early on for the linear-time solvable case $k = 2$: independently, Chvátal and Reed [7], Fernandez de la Vega [12], and Goerdt [17] proved $r_2 = 1$. For $k \geq 3$, neither the value nor the existence of r_k have been established. Friedgut [14], though, has proved the existence of a critical *sequence* $r_k(n)$ around which the probability of satisfiability goes from 1 to 0. In the following, we will take the liberty of writing $r_k \geq r^*$ if $F_k(n, rn)$ is satisfiable w.h.p. for all $r < r^*$ (and analogously for $r_k \leq r^*$).

Chvátal and Reed [7], besides proving $r_2 = 1$, gave the first lower bound for r_k , strengthening the positive-probability result of [6]. In particular, they considered a generalization of UC, called SC, which in the absence of unit clauses satisfies a random literal in a random 2-clause (and in the absence of 2-clauses satisfies a uniformly random literal). They proved that for all $k \geq 3$, if $r < (3/8)2^k/k$ then SC finds a satisfying truth assignment w.h.p.

In the last ten years, the satisfiability threshold conjecture has received attention in theoretical computer science, mathematics and, more recently, statistical physics. A large fraction of this attention has been devoted to the first computationally non-trivial case, $k = 3$ and a long series of results [4, 5, 16, 1, 3, 21, 10, 22, 19, 23, 20, 11, 13] has narrowed the potential range of r_3 . Currently this is pinned between 3.42 by Kaporis, Kirousis and Lalas [21] and 4.506 by Dubois and Boufkhad [10]. All upper bounds for r_3 come from probabilistic counting arguments, refining the idea of counting the expected number of satisfying truth assignments. All lower bounds on the other hand have been algorithmic, the refinement lying in considering progressively more sophisticated algorithms.

Unfortunately, for general k , neither of these two approaches above has helped narrow the asymptotic gap between the upper and lower bounds for r_k . The known techniques improve upon $r_k \leq 2^k \ln 2$ by a small additive constant, while the best lower bound, comes from Frieze and Suen's [16] analysis of a full generalization of UC:

Satisfy a random literal in a random shortest clause.

This gives $r_k \geq c_k 2^k/k$ where $\lim_{k \rightarrow \infty} c_k = 1.817 \dots$

If one chooses to live unencumbered by the burden of mathematical proof, then a powerful non-rigorous technique of statistical physics known as the “replica trick” is available. So far, predictions based on the replica trick have exhibited a strong (but not perfect) correlation with the (empirically observed) truth. Using this technique, Monasson and Zecchina [25] predicted $r_k \simeq 2^k \ln 2$. Like most arguments based on the replica trick, their argument is mathematically sophisticated but far from being rigorous.

If one indeed believes that the correct answer lies closer to the upper bound (for whatever reason) then analyzing more sophisticated satisfiability algorithms is an available option. Unfortunately, after a few steps down this path one is usually forced to choose between rather naive algorithms, which can be analyzed, or more sophisticated algorithms that might get closer to the threshold, but are much harder to analyze. In particular, the lack of progress over $c 2^k/k \leq r_k \leq 2^k \ln 2$ in the last ten years suggests the possibility that no (naive) algorithm can significantly improve the lower bound. At the same time, it is clear that proving lower bounds by analyzing algorithms is doing “more than we need”: we not only get a proof that a satisfying assignment exists but an explicit procedure for finding one.

In this paper, we eliminate the asymptotic gap for r_k by using the “second moment” method. Employing such a non-constructive argument allows us to overcome the limitations of current algorithmic techniques or, at least, of our capacity to analyze them. At the same time, not pursuing some particular satisfying truth assignment affords us a first, blurry glimpse of the “geometry” of the set of satisfying truth assignments. Our main result is the following.

Theorem 1 *For all $k \geq 2$, $r_k > 2^{k-1} \ln 2 - d_k$, where $d_k \rightarrow (1 + \ln 2)/2$.*

As we will see shortly, a straightforward application of the second moment method to random k -SAT fails rather dramatically: if X denotes the number of satisfying truth assignments, then $\mathbf{E}[X^2] > (1 + \epsilon)^n \mathbf{E}[X]^2$ for any $r > 0$. To prove Theorem 1 it will be crucial to focus on those satisfying truth assignments *whose complement is also satisfying*. Observe that this is equivalent to interpreting $F_k(n, m)$ as an instance of Not All Equal (NAE) k -SAT, where a truth assignment is NAE-satisfying if every clause contains at least one satisfied literal *and* at least one unsatisfied literal.

Analogously to random k -SAT, it is trivial to show that if $r > 2^{k-1} \ln 2 - (\ln 2)/2$ then w.h.p. $F_k(n, m = rn)$ has no NAE-satisfying truth assignments since their expected number is $o(1)$. We match this within an additive constant.

Theorem 2 *There exists a sequence $t_k \rightarrow 1/2$ such that if*

$$r < 2^{k-1} \ln 2 - (\ln 2)/2 - t_k$$

then w.h.p. $F_k(n, rn)$ is NAE-satisfiable.

Theorem 1 follows trivially from Theorem 2 since any NAE-satisfying assignment is also a satisfying assignment. Our method actually yields an explicit lower bound for the random NAE k -SAT threshold for each value of k as the solution to a transcendental equation (yet one without an attractive closed form, hence Theorem 2). It is, perhaps, worth comparing our lower bound for the NAE k -SAT threshold with the upper bound derived using the technique of [23] for small values of k . Even for $k = 3$, our lower bound is competitive with the best known lower bound of 1.514, obtained by analyzing a generalization of UC that minimizes the number of unit clauses [2]. For larger k , the gap between the upper and the lower bound rapidly converges to $\approx 1/4$.

| k | 3 | 5 | 7 | 10 | 12 |
|-------|-------|--------|--------|---------|----------|
| Lower | 3/2 | 9.973 | 43.432 | 354.027 | 1418.712 |
| Upper | 2.214 | 10.505 | 43.768 | 354.295 | 1418.969 |

Table 1. Bounds for the random NAE k -SAT threshold.

Recently, and independently of our work, Frieze and Wormald [15] showed that another way to successfully apply the second moment method to random k -SAT is to let k grow with n . In particular, let $\omega = k - \log_2 n \rightarrow \infty$, let $m_0 = -\frac{n \ln 2}{\ln(1-2^{-k})} = (2^k + O(1))n \ln 2$ and let $\epsilon = \epsilon(n) > 0$ be such that $\epsilon n \rightarrow \infty$. Then, $F_k(n, m)$ is w.h.p. satisfiable if $m < (1 - \epsilon)m_0$ but w.h.p. unsatisfiable if $m > (1 + \epsilon)m_0$.

We prove Theorem 2 by applying the following version of the second moment method (see Exercise 3.6 in [26]).

Lemma 1 *For any non-negative random variable X ,*

$$\Pr[X > 0] \geq \frac{\mathbf{E}[X]^2}{\mathbf{E}[X^2]} . \quad (1)$$

In particular, let $X \geq 0$ be the number of NAE-satisfying assignments of $F_k(n, m = rn)$. We will prove that for all $\epsilon > 0$ and all $k \geq k_0(\epsilon)$, if $r \leq 2^{k-1} \ln 2 - \frac{\ln 2}{2} - \frac{1}{2} - \epsilon$ then there exists some constant $C = C(k)$ such that

$$\mathbf{E}[X^2] < C \times \mathbf{E}[X]^2 .$$

By Lemma 1, this implies

$$\Pr[X > 0] \geq \Pr[F_k(n, rn) \text{ is NAE-satisfiable}] > 1/C .$$

To get Theorems 1 and 2 we boost this positive probability to $1 - o(1)$ by employing the following corollary of the aforementioned non-uniform threshold for random k -SAT [14] (and its analogue for random NAE k -SAT):

Corollary 1 *If $\liminf_{n \rightarrow \infty} \Pr[F_k(n, r^*n) \text{ is satisfiable}] > 0$, then $F_k(n, rn)$ is satisfiable w.h.p. for $r < r^*$.*

In the next section we give some intuition on why the second moment method fails when X is the number of satisfying truth assignments, and how letting X be the number of NAE-satisfying assignments rectifies the problem. In Section 3 we give some related general observations and point out potential connections to statistical physics. We lay the groundwork for bounding $\mathbf{E}[X^2]$ in Section 4. The actual bounding happens in Section 5. We conclude with some discussion in Section 6.

2. The second moment method


2.1. Random k -SAT

Let X denote the number of satisfying assignments of $F_k(n, m)$. Since X is the sum of 2^n indicator random variables, linearity of expectation implies that to bound $\mathbf{E}[X^2]$ we can consider all 4^n ordered pairs of truth assignments and bound the probability that both assignments in each pair are satisfying. It is easy to see that, by symmetry, for any pair of truth assignments s, t this probability depends only on the number of variables assigned the same value by s and t , i.e., their overlap. Thus, we can write $\mathbf{E}[X^2]$ as a sum with $n + 1$ terms, one for each possible value of the overlap z , the z th such term being: 2^n (counting over s) \times an “entropic” $\binom{n}{z}$ factor (counting overlap locations) \times a “correlation” factor measuring the probability that truth assignments s, t having overlap z are both satisfying.


Now, as we saw earlier, $\mathbf{E}[X] = [2(1 - 2^{-k})r]^n = c^n$. Thus, if r is such that $c < 1$, then $\Pr[X > 0] \leq \mathbf{E}[X] = o(1)$ and we readily know that $F_k(n, rn)$ is w.h.p. unsatisfiable. (Note that $\Pr[X > 0] = o(1)$ even when $c = 1$ since the naive upper bound is not tight.) Therefore, we are only interested in the case where $\mathbf{E}[X^2] \geq \mathbf{E}[X]^2 = (1 + \epsilon)^n$ for some $\epsilon = \epsilon(r) > 0$. Since the sum defining $\mathbf{E}[X^2]$ has only $n + 1$ terms we see that, up to polynomial factors, $\mathbf{E}[X^2]$ is equal to the contribution of the term maximizing the “entropy-correlation” product.

Observe, now, that if $z = n/2$, then the probability that s and t are both satisfying is the square of the probability that one of them is. To see this take s to be, say, the all 0s assignment and consider the set of clauses this precludes from being in the formula. Thus, for truth assignments that overlap on $n/2$ bits, the events of being satisfying are independent. Therefore, up to polynomial factors, $\mathbf{E}[X^2]$ is equal to the $z = n/2$ term of the sum defining $\mathbf{E}[X^2]$.

From the above discussion, letting $\alpha = z/n$, we see that if the entropy-correlation factor is maximized at some $\alpha \neq 1/2$ then the second moment method fails. On the other hand, as we will see, if the maximum does indeed occur at $\alpha = 1/2$, then the polynomial factors cancel out and the ratio $\mathbf{E}[X^2]/\mathbf{E}[X]^2$ is bounded by a constant independent of n , implying that in that case $\Pr[X > 0] > 1/C$.



$k = 5, r = 16, 18, 20, 22, 24$ (top to bottom)



$k = 5, r = 8, 9, 10, 11, 12$ (top to bottom)

Figure 1. The n th root of $2 \times$ the entropy-correlation product for k -SAT as a function of the overlap $\alpha = z/n$.

With these observations in mind, in Fig. 1 we plot the n th root of each of the $n+1$ terms contributing to $\mathbf{E}[X^2]$ as a function of $\alpha = z/n$ for $k = 5$ and different values of r .


Unfortunately, we see that for all values of r considered the maximum lies to the right of $\alpha = 1/2$. The reason for this is that the correlation factor for k -SAT is strictly increasing with $\alpha = z/n$. For instance, as we saw above, if s is satisfying and t has an overlap of $z = n/2$ with s , then the conditional probability that t is also satisfying equals its *a priori* value $(1 - 1/2^k)^m$. But if z decreases, say, to 0 then the conditional probability that t is satisfying decreases to $(1 - 1/(2^k - 1))^m$, penalizing $t = \bar{s}$ exponentially and making it the least likely assignment to be satisfying.

This asymmetry in the correlation factor implies that for all $r > 0$ its product with the (symmetric) entropy factor is maximized at some $\alpha > 1/2$. Therefore, $\mathbf{E}[X^2]$ is greater than $\mathbf{E}[X]^2$ by an exponential factor for all $r > 0$, and Lemma 1 fails to give any non-trivial lower bound. To have any hope of getting a lower bound by the second method we need to consider a set of satisfying assignments for which the derivative of the correlation factor at $1/2$ is zero.

2.2. Random NAE k -SAT

One attractive feature of the second moment method is that we are free to apply it to any random variable X such that $X > 0$ implies that $F_k(n, m)$ is satisfiable. In particular, we can refine our earlier application of the method by focusing on any subset of the set of satisfying assignments.

Considering only assignments that are NAE-satisfying — or, equivalently, whose complement is also satisfying — makes the correlation factor symmetric around $\alpha = 1/2$ as twin satisfying assignments s and \bar{s} provide an equal “tug” to every other truth assignment t . As a result, we always have a local extremum at $\alpha = 1/2$ since both the correlation factor and the entropy are symmetric around it. Moreover, since the entropic term is independent for r , this extremum is a local maximum for sufficiently small r . Whenever this is also the *global* maximum, the second moment succeeds.



$k = 5, r = 9.973$

Figure 2. The n th root of $2 \times$ the entropy-correlation product for NAE k -SAT as a function of the overlap $\alpha = z/n$.

In Fig. 2 we plot the n th root of the entropy-correlation product for NAE k -SAT for various values of r . Let us start with the top picture, where $k = 5$ and r increases from 8 to 12 as we go from top to bottom. For $r = 8, 9$ we see that, indeed, the global maximum occurs at $\alpha = 1/2$. As result, for such r we have $\mathbf{E}[X^2] = \Theta(\mathbf{E}[X]^2)$, implying that the formula is NAE-satisfiable with positive probability.

For the cases $r = 11, 12$, on the other hand, we see that at $\alpha = 1/2$ the function has dropped below 1 and therefore $\mathbf{E}[X]^2 = o(1)$, implying that w.h.p. $F_k(n, rn)$ has no NAE-satisfying truth assignment. It is worth noting that for $r = 11$ we have $\Pr[X > 0] = o(1)$, even though $\mathbf{E}[X^2]$ is exponentially large, due to the maxima close to 0 and 1.

The most interesting case is $r = 10$ where $\alpha = 1/2$ is a local maximum (and greater than 1) but the global maxima occur at 0.08, 0.92 where the function equals 1.0145... (vs. 1.0023... at $\alpha = 1/2$). Because of this, we have $\mathbf{E}[X^2]/\mathbf{E}[X]^2 > (1.0144/1.0024)^n$, implying that the second moment method only gives an exponentially small lower bound on $\Pr[X > 0]$ in spite of the fact that the expected number of NAE-satisfying truth assignments is exponential. Note, also, that according to Table 1 the best known upper bound for $k = 5$ is $10.505 > 10$.

Indeed, the largest value for which the second moment succeeds for $k = 5$ is $r = 9.973...$ This is depicted in the bottom picture where the three peaks have the same height. For $r > 9.973$ the peaks near 0 and 1 surpass the one at $\alpha = 1/2$ and the second moment method fails.

3. Intuition

3.1. Reducing the variance

Given two truth assignments s, t that have overlap z let

$$\text{boost}(z) = \frac{\Pr[t \text{ is satisfying} \mid s \text{ is satisfying}]}{\Pr[t \text{ is satisfying}]} .$$

It is not hard to see that

$$\frac{\mathbf{E}[X^2]}{\mathbf{E}[X]^2} = 2^{-n} \sum_{z=0}^n \binom{n}{z} \text{boost}(z) .$$

To examine one particular source contributing to $\text{boost}(z)$ in the case of random k -SAT, it is helpful to introduce the following quantity: given a truth assignment s and a formula F let $Q = Q(s, F)$ be the total number of literal occurrences in F that are satisfied by s . Thus, $Q(s, F)$ is maximized when s assigns each variable its “majority” value.

It is well-known that, with respect to properties that hold w.h.p., $F_k(n, m = rn)$ is equivalent to a random formula generated as follows: first, for each literal ℓ , generate R_ℓ literal occurrences, where the $\{R_\ell\}$ are i.i.d. Poisson random variables with mean $kr/2$; then, partition these literal occurrences randomly into m parts of size k .

In this model we can easily factor $\Pr[s \text{ is satisfying}]$ as $\sum_q \Pr[Q = q] \times \Pr[s \text{ is satisfying} \mid Q = q]$. Clearly, the probability of Q deviating significantly from its expected value $km/2$ is exponentially small. At the same time, though, any such increase in Q affords s tremendous advantage in terms of its likelihood to be satisfying. Moreover, since w.h.p. each variable appears in $O(\log n)$ clauses, this advantage will be very much shared with the truth assignments having large overlap with s , thus contributing heavily to the boost function and, as a result, to $\mathbf{E}[X^2]$.

On the other hand, if we consider the probability that s is NAE-satisfying it is clear that s would like Q to be as close as possible to $km/2$. In other words, now the typical case is the most favorable case and the clustering around truth assignments that satisfy many literal occurrences disappears.

Whether this is the main reason for which the second moment method succeeds for random NAE k -SAT remains an interesting question. Considering *regular* random k -SAT, where all literals are required to appear an equal number of times, seems like an interesting test of this hypothesis.

3.2. Geometry and connections to statistical physics

Statistical physicists have developed a number of methods for investigating phase transitions which, while non-rigorous, are often in spectacular agreement with numerical and experimental results. One of these methods is the replica trick. The term “replica” comes from the fact that

when q is an integer one can compute $\mathbf{E}[X^q]$ by considering the interactions between q elementary objects, or “replicas”, counted by X . In our case, we consider two truth assignments when calculating the second moment.

At a high level, the replica trick amounts to computing $\mathbf{E}[\ln X]$ by calculating $\mathbf{E}[X^q]$ for all *integer* q and then plugging in the resulting formula to the expression $\mathbf{E}[\ln X] = \lim_{q \rightarrow 0} (\mathbf{E}[X^q] - 1)/q$. The fundamental leap of faith, of course, lies in allowing the analytic continuation $q \rightarrow 0$ from integer values of q . Even to get this far, however, one has to deal with the often daunting task of computing $\mathbf{E}[X^q]$ for all integer q .

When X counts objects expressed as binary strings, such as satisfying assignments, to calculate $\mathbf{E}[X^q]$ one must in general maximize a function of $2^q - 1$ “overlaps”, each overlap counting the number of variables assigned a given q -vector of 0/1 values by the q assignments/replicas. (Note that in random [NAE] k -SAT, since variables are negated randomly in each clause, we can take one of the q assignments to be the all 0s, so we only have $2^{q-1} - 1$ overlaps.)

By taking another leap of faith, one can dramatically reduce the dimensionality of this maximization problem to q by assuming *replica symmetry*, i.e., that the global maximum is symmetric under permutations of the replicas. For satisfiability problems this means that all overlap variables with the same number of 1s in their respective q -vector take the same value. While this assumption is often wrong, it can lead to good approximations. In particular, replica symmetry was assumed in the work of Monasson and Zecchina [25] predicting $r_k \simeq 2^k \ln 2$.

A standard indicator of the plausibility of replica symmetry in a given system is the (usually experimentally measured) distribution of overlaps between randomly chosen ground states, in our case satisfying assignments. If replica symmetry holds, this distribution is tightly peaked around its mean; if not, i.e., if “replica symmetry breaking” takes place, this distribution typically gains multiple peaks or becomes continuous in some open interval.

Intriguingly, the second-moment method is essentially a calculation of the overlap distribution in the *annealed approximation*, i.e., after we average over random formulas (giving formulas with more satisfying assignments a heavier influence in the overlap distribution). For random NAE k -SAT we saw that, almost all the way to the threshold, the overlap distribution is sharply concentrated around $n/2$, since when we take n th powers the contribution of all other terms vanishes.

In other words, we have shown that in the annealed approximation, the overlap distribution behaves *as if the NAE-satisfying assignments were scattered independently throughout the hypercube*.

4. Groundwork

Let X be the number of NAE-satisfying assignments of $F_k(n, m = rn)$. We start by calculating $\mathbf{E}[X]$. For any given assignment s , the probability that a random clause is satisfied by s is the probability that its k literals are neither all true nor all false. We call this probability $p = 1 - 2^{1-k}$.

Since clauses are drawn independently with replacement and we have $m = rn$ clauses, we see that

$$\mathbf{E}[X] = (2p^r)^n. \quad (2)$$

To calculate $\mathbf{E}[X^2]$ we first observe that, by linearity of expectation, it is equal to the expected number of ordered pairs of truth assignments s, t such that both s and t are NAE-satisfying. We claim that the probability that a pair of truth assignments s, t are both NAE-satisfying depends only on the number of variables to which they assign the same value (their overlap). In particular, we claim that if s and t have overlap $z = \alpha n$, where $0 \leq \alpha \leq 1$, then a random k -clause c is satisfied by both s and t with probability

$$\begin{aligned} f(\alpha) &= 1 - 2 \cdot 2^{1-k} + 2^{1-k}(\alpha^k + (1-\alpha)^k) \\ &= 1 - 2^{1-k}(2 - \alpha^k - (1-\alpha)^k). \end{aligned} \quad (3)$$

To see this, first recall that the probability of a clause c not being satisfied by s is $1 - p = 2^{1-k}$. Moreover, if c is not satisfied by s , then in order for c to also not be satisfied by t , it must be that either all the variables in c have the same value in t and s , or they all have opposite values. Since s and t have an overlap of $z = \alpha n$ variables and the variables in each clause are distinct, the probability of this last event is $\alpha^k + (1-\alpha)^k$. Thus, (3) follows by inclusion-exclusion.

Now, since the number of ordered pairs of assignments with overlap z is $2^n \binom{n}{z}$ and since the $m = rn$ clauses are drawn independently and with replacement we see that

$$\mathbf{E}[X^2] = 2^n \sum_{z=0}^n \binom{n}{z} f(z/n)^{rn}. \quad (4)$$

We will bound this sum by focusing on its largest terms. The proof of the following lemma, based on standard asymptotic techniques, appears in the Appendix.

Lemma 2 *Let F be a real analytic positive function on $[0, 1]$ and define g on $[0, 1]$ as*

$$g(\alpha) = \frac{F(\alpha)}{\alpha^\alpha (1-\alpha)^{1-\alpha}},$$

where $0^0 \equiv 1$. *If there exists $\alpha_{\max} \in (0, 1)$ such that $g(\alpha_{\max}) \equiv g_{\max} > g(\alpha)$ for all $\alpha \neq \alpha_{\max}$, and $g''(\alpha_{\max}) < 0$, then there exist constants $B, C > 0$ such that for all sufficiently large n*

$$B \times g_{\max}^n \leq \sum_{z=0}^n \binom{n}{z} F(z/n)^n \leq C \times g_{\max}^n.$$

With Lemma 2 in mind we define

$$g_r(\alpha) = \frac{f(\alpha)^r}{\alpha^\alpha (1-\alpha)^{1-\alpha}}.$$

We will prove that

Lemma 3 *For every $\epsilon > 0$, there exists $k_0 = k_0(\epsilon)$ such that for all $k \geq k_0$, if*

$$r \leq 2^{k-1} \ln 2 - \frac{\ln 2}{2} - \frac{1}{2} - \epsilon$$

then $g_r(\alpha) < g_r(1/2)$ for all $\alpha \neq 1/2$, and $g_r''(1/2) < 0$.

Therefore, for all r, k, ϵ as in Lemma 3,

$$\mathbf{E}[X^2] < C \times (2g_r(1/2))^n,$$

where $C = C(k)$ is independent of n . At the same time, observe that $\mathbf{E}[X]^2 = (4p^{2r})^n = (2g_r(1/2))^n$. Therefore, for all r, k, ϵ as in Lemma 3,

$$\frac{\mathbf{E}[X^2]}{\mathbf{E}[X]^2} < C,$$

which, by Lemma 1, implies

$$\Pr[X > 0] > 1/C.$$

Thus, along with Corollary 1, Lemma 3 suffices to establish Theorems 1 and 2.

5. Proof of Lemma 3

We wish to show that $g_r''(1/2) < 0$ and that $g_r(\alpha) < g_r(1/2)$ for all $\alpha \neq 1/2$. Since g_r is symmetric around $1/2$, we can restrict to $\alpha \in (1/2, 1]$. We will divide this interval into two parts and handle them with two separate lemmata. The first lemma deals with $\alpha \in (1/2, 0.9]$ and also establishes that $g_r''(1/2) < 0$.

Lemma 4 *Let $\alpha \in (1/2, 0.9]$. For all $k \geq 74$, if $r \leq 2^{k-1} \ln 2$ then $g_r(\alpha) < g_r(1/2)$ and $g_r''(1/2) < 0$.*

The second lemma deals with $\alpha \in (0.9, 1]$.

Lemma 5 *Let $\alpha \in (0.9, 1]$. For every $\epsilon > 0$ and all $k \geq k_0(\epsilon)$, if $r \leq 2^{k-1} \ln 2 - \frac{\ln 2}{2} - \frac{1}{2} - \epsilon$ then $g_r(\alpha) < g_r(1/2)$.*

Combining Lemmata 4 and 5 we see that for every $\epsilon > 0$, there exists $k_0 = k_0(\epsilon)$ such that for all $k \geq k_0$ if

$$r \leq 2^{k-1} \ln 2 - \frac{\ln 2}{2} - \frac{1}{2} - \epsilon$$

then $g_r(\alpha) < g_r(1/2)$ for all $\alpha \neq 1/2$ and $g_r''(1/2) < 0$, establishing Lemma 3. We prove Lemmata 4 and 5 below.

The reader should keep in mind that we have made no attempt to optimize the value of k_0 in Lemma 5, opting instead for proof simplicity.

Proof of Lemma 4. We will first prove that for $k \geq 74$, g_r is strictly decreasing in $\alpha = (1/2, 0.9]$, thus establishing $g_r(\alpha) < g_r(1/2)$. Since g_r is positive, to do this it suffices to prove that $(\ln g_r)' = g'_r/g_r < 0$ in this interval. In fact, since $g'_r(\alpha) = (\ln g_r)' = 0$ at $\alpha = 1/2$, it will suffice to prove that for $\alpha \in [1/2, 0.9]$ we have $(\ln g_r)'' < 0$. Now,

$$\begin{aligned} (\ln g_r(\alpha))'' &= r \left(\frac{f''(\alpha)}{f(\alpha)} - \frac{f'(\alpha)^2}{f(\alpha)^2} \right) - \frac{1}{\alpha(1-\alpha)} \\ &\leq r \frac{f''(\alpha)}{f(\alpha)} - \frac{1}{\alpha(1-\alpha)}. \end{aligned} \quad (5)$$

To show that the r.h.s. of (5) is negative we first note that for $\alpha \geq 1/2$ and $k > 3$,

$$f''(\alpha) = 2^{1-k} k(k-1)(\alpha^{k-2} + (1-\alpha)^{k-2}) < 2^{2-k} \alpha^{k-2} k^2$$

is monotonically increasing. Therefore,

$$f''(\alpha) \leq f''(0.9) < 2^{2-k} 0.9^{k-2} k^2.$$

Moreover, for all α , $f(\alpha) \geq f(1/2) = (1 - 2^{-k})^2$. Therefore, since $1/(\alpha(1-\alpha)) \geq 4$ and $r \leq 2^{k-1} \ln 2$, it suffices to observe that for all $k \geq 74$,

$$(2^{k-1} \ln 2) \times \frac{2^{2-k} 0.9^{k-2} k^2}{(1 - 2^{-74})^2} - 4 < 0.$$

Finally, recalling that $g'(1/2) = 0$ and using

$$(\ln g_r)'' = \frac{g''(\alpha)}{g(\alpha)} - \frac{g'(\alpha)^2}{g(\alpha)^2}$$

we see that at $g''_r(1/2) < 0$ since $(\ln g_r)''(1/2) < 0$.

Proof of Lemma 5. By the definition of g_r we see that $g_r(\alpha) < g_r(1/2)$ if and only if

$$\left(\frac{f(\alpha)}{f(1/2)} \right)^r < 2\alpha^\alpha (1-\alpha)^{1-\alpha}. \quad (6)$$

Letting $h(\alpha) = -\alpha \ln \alpha - (1-\alpha) \ln(1-\alpha)$ denote the entropy function, we see that (6) holds as long as

$$\frac{r}{\ln 2 - h(\alpha)} < \frac{1}{\ln(1+w)}$$

where

$$w = \frac{f(\alpha) - f(1/2)}{f(1/2)}.$$

Observe now that for $k > 3$, f is strictly increasing in $(1/2, 1]$, so $w > 0$. Moreover, for any $x > 0$

$$\frac{1}{\ln(1+x)} \geq \frac{1}{x} + \frac{1}{2} - \frac{x}{12}.$$

Since $f(\alpha) - f(1/2) = 2^{1-k}(\alpha^k + (1-\alpha)^k - 2^{1-k}) < 2^{1-k}$ and $f(1/2) = (1 - 2^{1-k})^2 > 1 - 2^{2-k}$, we thus see that it suffices to have

$$\frac{r}{\ln 2 - h(\alpha)} < \frac{2^{k-1} - 2}{\alpha^k + (1-\alpha)^k - 2^{1-k}} + \frac{1}{2} - \frac{2^{1-k}}{12}. \quad (7)$$

Now observe that for any $0 < \alpha < 1$ and $0 \leq q < \alpha^k$,

$$\frac{1}{\alpha^k - q} \geq 1 + k(1-\alpha) + q.$$

Since $\alpha > 1/2$ we can set $q = 2^{1-k} - (1-\alpha)^k$, yielding

$$\frac{1}{\alpha^k + (1-\alpha)^k - 2^{1-k}} \geq 1 + k(1-\alpha) + 2^{1-k} - (1-\alpha)^k.$$

Since $2^k(1-\alpha)^k < 5^{-k}$, we find that (6) holds as long as $r \leq \phi(y) - 2^{-k}$ where

$$\phi(\alpha) \equiv (\ln 2 - h(\alpha)) \left(2^{k-1} + (2^{k-1} - 2)k(1-\alpha) - \frac{1}{2} \right).$$

We are thus left to minimize ϕ in $(0.9, 1]$. Since ϕ is analytic its minima can only occur at 0.9 or 1, or where $\phi' = 0$. The derivative of ϕ is

$$\begin{aligned} \phi'(\alpha) &= (2^{k-1} - 2) \times \left[-k(\ln 2 - h(\alpha)) \right. \\ &\quad \left. + (\ln \alpha - \ln(1-\alpha)) \left(1 + k(1-\alpha) + \frac{3}{2^k - 4} \right) \right]. \end{aligned} \quad (8)$$

Note now that for all $k > 1$

$$\lim_{\alpha \rightarrow 1} \phi'(\alpha) = -\frac{2^k - 1}{2} \ln(1-\alpha)$$

is positively infinite. At the same time,

$$\phi'(0.9) < -0.07 \times 2^k k + 1.1(2^k - 1) + 0.3k$$

is negative for $k \geq 16$. Therefore, ϕ is minimized in the interior of $(0.9, 1]$ for all $k \geq 16$. Setting ϕ' to zero gives

$$-\ln(1-\alpha) = \frac{k(\ln 2 - h(\alpha))}{1 + k(1-\alpha) + 3/(2^k - 4)} - \ln \alpha. \quad (9)$$

By “bootstrapping” we derive a tightening series of lower bounds on the solution for the l.h.s. of (9) for $\alpha \in (0.9, 1)$. Note first that we have an easy upper bound,

$$-\ln(1-\alpha) < k \ln 2 - \ln \alpha. \quad (10)$$

At the same time, if $k > 2$ then $3/(2^k - 4) < 1$, implying

$$-\ln(1-\alpha) > \frac{k(\ln 2 - h(\alpha))}{2 + k(1-\alpha)} - \ln \alpha. \quad (11)$$

If we write $k(1 - \alpha) = B$ then (11) becomes

$$-\ln(1 - \alpha) > \frac{\ln 2 - h(\alpha)}{1 - \alpha} \left(\frac{B}{B + 2} \right) - \ln \alpha . \quad (12)$$

By inspection, if $B \geq 3$ the r.h.s. of (12) is greater than the l.h.s. for all $\alpha > 0.9$, yielding a contradiction. Therefore, $k(1 - \alpha) < 3$ for all $k > 2$. Since $\ln 2 - h(\alpha) > 0.36$ for $\alpha > 0.9$, we see that for $k > 2$, (11) implies

$$-\ln(1 - \alpha) > 0.07 k . \quad (13)$$

Observe now that, by (13), $k(1 - \alpha) < k e^{-0.07k}$ and, hence, as k increases the denominator of (9) approaches 1.

To bootstrap, we note that since $\alpha > 1/2$ we have

$$h(\alpha) \leq -2(1 - \alpha) \ln(1 - \alpha) \quad (14)$$

$$\begin{aligned} &< 2e^{-0.07k} (k \ln 2 - \ln 0.9) \\ &< 2k e^{-0.07k} \end{aligned} \quad (15)$$

where (15) relies on (10), (13). Moreover, $\alpha > 1/2$ implies $-\ln \alpha \leq 2(1 - \alpha) < 2e^{-0.07k}$. Thus, by using (13) and the fact $1/(1+x) > 1-x$ for all $x > 0$, (9) gives for $k \geq 3$,

$$\begin{aligned} -\ln(1 - \alpha) &> \frac{k(\ln 2 - h(\alpha))}{1 + k(1 - \alpha) + 3/(2^k - 4)} \\ &> \frac{k(\ln 2 - 2k e^{-0.07k})}{1 + 2k e^{-0.07k}} \\ &> k(\ln 2 - 2k e^{-0.07k})(1 - 2k e^{-0.07k}) \\ &> k \ln 2 - 4k^2 e^{-0.07k} . \end{aligned} \quad (16)$$

For $k \geq 166$, $4k^2 e^{-0.07k} < 1$. Thus, by (16), we have $1 - \alpha < 3 \times 2^{-k}$. This, in turn, implies $-\ln \alpha \leq 2(1 - \alpha) < 6 \times 2^{-k}$ and so, by (14) and (10), we have for $\alpha > 0.9$

$$h(\alpha) < 6 \times 2^{-k} (k \ln 2 - \ln \alpha) < 5k 2^{-k} . \quad (17)$$

Plugging (17) into (9) to bootstrap again, we get that for $k \geq 3$

$$\begin{aligned} -\ln(1 - \alpha) &> \frac{k(\ln 2 - 5k 2^{-k})}{1 + 3k 2^{-k} + 3/(2^k - 4)} \\ &> \frac{k(\ln 2 - 5k 2^{-k})}{1 + 6k 2^{-k}} \\ &> k(\ln 2 - 5k 2^{-k})(1 - 6k 2^{-k}) \\ &> k \ln 2 - 11k^2 2^{-k} . \end{aligned}$$

Since $e^x < 1 + 2x$ for $x < 1$ and $11k^2 2^{-k} < 1$ for $k > 10$, we see that for such k

$$1 - \alpha < 2^{-k} + 22k^2 2^{-2k} .$$

Plugging into (10) the fact $-\ln \alpha < 6 \times 2^{-k}$ we get $-\ln(1 - \alpha) < k \ln 2 + 6 \times 2^{-k}$. Using that $e^{-x} \geq 1 - x$ for $x \geq 0$, we get the closely matching upper bound,

$$1 - \alpha > 2^{-k} - 6 \times 2^{-2k} .$$

Thus, we see that for $k \geq 166$, ϕ is minimized at an α_{\min} which is within δ of $1 - 2^{-k}$, where $\delta = 22k^2 2^{-2k}$. Let T be the interval $[1 - 2^{-k} - \delta, 1 - 2^{-k} + \delta]$. Clearly the minimum of ϕ is at least $\phi(1 - 2^{-k}) - \delta \times \max_{\alpha \in T} |\phi'(\alpha)|$. It is easy to see from (8) that if $\alpha \in T$ then $|\phi'(\alpha)| \leq 2k 2^k$.

Now, a simple calculation using that $\ln(1 - 2^{-k}) > -2^{-k} - 2^{-2k}$ for $k \geq 1$ gives

$$\begin{aligned} \phi(1 - 2^{-k}) &= \frac{1}{2} ((2^k - k) \ln 2 + (2^k - 1) \ln(1 - 2^{-k})) \\ &\quad \times (1 + (k - 1) 2^{-k} - k 2^{2-2k}) \\ &> 2^{k-1} \ln 2 - \frac{\ln 2}{2} - \frac{1}{2} - k^2 2^{-k} . \end{aligned}$$

Therefore,

$$\phi_{\min} \geq 2^{k-1} \ln 2 - \frac{\ln 2}{2} - \frac{1}{2} - 45k^3 2^{-k} .$$

Finally, recall that (6) holds as long as $r < \phi_{\min} - 2^{-k}$, i.e.,

$$r < 2^{k-1} \ln 2 - \frac{\ln 2}{2} - \frac{1}{2} - 46k^3 2^{-k} .$$

Clearly, we can take $k_0 = O(\ln \epsilon^{-1})$ so that for all $k \geq k_0$ the error term $46k^3 2^{-k}$ is smaller than any $\epsilon > 0$.

6. Conclusions

We have shown that the second moment method can be used to determine the random k -SAT threshold within a factor of 2. We also showed that it gives extraordinarily tight bounds for random NAE k -SAT, determining the threshold for that problem within a small additive constant.

At this point, it seems vital to understand the following:

1. Why does the second moment method perform so well for NAE k -SAT? The symmetry of this problem explains why the method gives a non-trivial bound, but not why it gives essentially the exact answer.
2. How can we close the factor of 2 gap for the random k -SAT threshold? Are there other large subsets of satisfying assignments that are not strongly correlated?
3. Does the geometry of the set of satisfying assignments have any implications for algorithms? Perhaps more modestly(?), is there a polynomial-time algorithm that succeeds with positive probability for $r = \omega(k) 2^k/k$, where $\omega(k) \rightarrow \infty$? What about $\omega(k) = \Theta(k)$?

Acknowledgements We would like to thank Henry Cohn for bringing [9] to our attention and Chris Umans for helpful comments regarding the presentation. Special thanks to Paul Beame for reading earlier versions and making many helpful suggestions.

References

- [1] D. Achlioptas. Setting two variables at a time yields a new lower bound for random 3-SAT. In *Proc. 32nd Annual ACM Symposium on Theory of Computing*, pages 28–37, 2000.
- [2] D. Achlioptas, A. Chtcherba, G. Istrate, and C. Moore. The phase transition in 1-in- k SAT and NAE 3-SAT. In *Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 721–722, 2001.
- [3] D. Achlioptas and G. B. Sorkin. Optimal myopic algorithms for random 3-SAT. In *Proc. 41st Annual Symposium on Foundations of Computer Science*, pages 590–600, 2000.
- [4] A. Z. Broder, A. M. Frieze, and E. Upfal. On the satisfiability and maximum satisfiability of random 3-CNF formulas. In *Proc. 4th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 322–330, 1993.
- [5] M.-T. Chao and J. Franco. Probabilistic analysis of two heuristics for the 3-satisfiability problem. *SIAM J. Comput.*, 15(4):1106–1118, 1986.
- [6] M.-T. Chao and J. Franco. Probabilistic analysis of a generalization of the unit-clause literal selection heuristics for the k -satisfiability problem. *Inform. Sci.*, 51(3):289–314, 1990.
- [7] V. Chvátal and B. Reed. Mick gets some (the odds are on his side). In *Proc. 33th Annual Symposium on Foundations of Computer Science*, pages 620–627, 1992.
- [8] V. Chvátal and E. Szemerédi. Many hard examples for resolution. *J. Assoc. Comput. Mach.*, 35(4):759–768, 1988.
- [9] N. G. de Bruijn. *Asymptotic methods in analysis*. Dover Publications Inc., New York, third edition, 1981.
- [10] O. Dubois, Y. Boufkhad, and J. Mandler. Typical random 3-SAT formulae and the satisfiability threshold. In *Proc. 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 126–127, 2000.
- [11] A. El Maftouhi and W. Fernandez de la Vega. On random 3-sat. *Combin. Probab. Comput.*, 4(3):189–195, 1995.
- [12] W. Fernandez de la Vega. On random 2-SAT. Manuscript, 1992.
- [13] J. Franco and M. Paull. Probabilistic analysis of the Davis–Putnam procedure for solving the satisfiability problem. *Discrete Appl. Math.*, 5(1):77–87, 1983.
- [14] E. Friedgut. Necessary and sufficient conditions for sharp thresholds of graph properties, and the k -SAT problem. *J. Amer. Math. Soc.*, 12:1017–1054, 1999.
- [15] A. Frieze and N. C. Wormald. Random k -SAT: a tight threshold for moderately growing k . In *Proceedings of the 5th International Symposium on Theory and Applications of Satisfiability Testing*, pages 1–6, 2002.
- [16] A. M. Frieze and S. Suen. Analysis of two simple heuristics on a random instance of k -SAT. *J. Algorithms*, 20(2):312–355, 1996.
- [17] A. Goerdt. A threshold for unsatisfiability. *J. Comput. System Sci.*, 53(3):469–486, 1996.
- [18] A. Haken. The intractability of resolution. *Theoret. Comput. Sci.*, 39(2-3):297–308, 1985.
- [19] S. Janson, Y. C. Stamatiou, and M. Vamvakari. Bounding the unsatisfiability threshold of random 3-SAT. *Random Structures Algorithms*, 17(2):103–116, 2000.
- [20] A. Kamath, R. Motwani, K. Palem, and P. Spirakis. Tail bounds for occupancy and the satisfiability threshold conjecture. *Random Structures Algorithms*, 7(1):59–80, 1995.
- [21] A. Kaporis, L. M. Kirousis, and E. Lalas. Personal communication. April, 2002.
- [22] A. Kaporis, L. M. Kirousis, Y. C. Stamatiou, M. Vamvakari, and M. Zito. The unsatisfiability threshold revisited. *Disc. Math.*, to appear.
- [23] L. M. Kirousis, E. Kranakis, D. Krizanc, and Y. Stamatiou. Approximating the unsatisfiability threshold of random formulas. *Random Structures Algorithms*, 12(3):253–269, 1998.
- [24] D. G. Mitchell, B. Selman, and H. J. Levesque. Hard and easy distributions of SAT problems. In *Proc. 10th National Conference on Artificial Intelligence*, pages 459–462, 1992.
- [25] R. Monasson and R. Zecchina. Statistical mechanics of the random K -satisfiability model. *Phys. Rev. E* (3), 56(2):1357–1370, 1997.
- [26] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, 1995.
- [27] A. Urquhart. Hard examples for resolution. *J. Assoc. Comput. Mach.*, 34(1):209–219, 1987.

A. Proof of Lemma 2

The idea is that because of the binomial coefficient, the sum only has $\Theta(\sqrt{n})$ “significant” terms, each of which is of size $\Theta(g_{\max}^n/\sqrt{n})$. The proof amounts to replacing the sum by an integral and then using the Laplace method for asymptotic integrals [9].

We prove the upper bound first. Recall the following form of Stirling’s approximation, valid for all $n > 0$:

$$\begin{aligned} n! &> \sqrt{2\pi n} n^n e^{-n} \left(1 + \frac{1}{12n}\right) \\ n! &< \sqrt{2\pi n} n^n e^{-n} \left(1 + \frac{1}{6n}\right). \end{aligned}$$

Thus, for any $0 < z \leq n/2$, letting $\alpha = z/n$ we have

$$\begin{aligned} \binom{n}{\alpha n} &< \frac{1}{\sqrt{2\pi n}} \frac{1}{\sqrt{\alpha(1-\alpha)}} \left(\frac{1}{\alpha^\alpha (1-\alpha)^{1-\alpha}}\right)^n \\ &\quad \times \frac{1 + 1/(6n)}{1 + 1/(12z)} \\ &\leq \frac{1}{\sqrt{2\pi n}} \frac{1}{\sqrt{\alpha(1-\alpha)}} \left(\frac{1}{\alpha^\alpha (1-\alpha)^{1-\alpha}}\right)^n \quad (18) \end{aligned}$$

and, similarly, for any $0 < z < n$ we have

$$\begin{aligned} \binom{n}{\alpha n} &> \frac{1}{\sqrt{2\pi n}} \frac{1}{\sqrt{\alpha(1-\alpha)}} \left(\frac{1}{\alpha^\alpha (1-\alpha)^{1-\alpha}}\right)^n \\ &\quad \times \frac{1}{(1 + 1/(6z))(1 + 1/(6(n-z)))} \\ &> \frac{36}{49} \frac{1}{\sqrt{2\pi n}} \frac{1}{\sqrt{\alpha(1-\alpha)}} \left(\frac{1}{\alpha^\alpha (1-\alpha)^{1-\alpha}}\right)^n. \end{aligned}$$

To prove the upper bound, we use (18) to write

$$\sum_{z=0}^n \binom{n}{z} F(z/n)^n \leq \frac{1}{\sqrt{2\pi n}} \sum_{0 < z < n} \frac{g(z/n)^n}{\sqrt{(z/n)(1-z/n)}} + F(0)^n + F(1)^n. \quad (19)$$

Let $\epsilon = \min\{\alpha_{\max}, 1 - \alpha_{\max}\}/2 > 0$. Let $g_* < g_{\max}$ be the maximum value of g in $[0, \epsilon] \cup [1 - \epsilon, 1]$. Since $g_* \geq g(0) = F(0)$ and $g_* \geq g(1) = F(1)$, using (19) we get

$$\begin{aligned} & \sum_{z=0}^n \binom{n}{z} F(z/n)^n \\ & \leq \frac{1}{\sqrt{2\pi n}} \times \sum_{z=\epsilon n}^{(1-\epsilon)n} \frac{g(z/n)^n}{\sqrt{(z/n)(1-z/n)}} + n^{3/2} g_*^n \\ & \leq \frac{1}{\sqrt{2\pi n}} \frac{1}{\sqrt{\epsilon(1-\epsilon)}} \times \sum_{z=\epsilon n}^{(1-\epsilon)n} g(z/n)^n + n^{3/2} g_*^n. \end{aligned} \quad (20)$$

Next, we wish to replace the sum in (20) with an integral. We first recall that for any integrable function ϕ that is monotone in $[a, b]$

$$\left| \sum_{j=0}^s \phi\left(a + j \frac{(b-a)}{s}\right) - \frac{s}{b-a} \int_a^b \phi(x) dx \right| \leq \max\{F(a), F(b)\}.$$

Therefore if ϕ has M extrema in $[a, b]$, we can divide $[a, b]$ into $M+1$ intervals on which ϕ is monotone, giving

$$\left| \sum_{j=0}^s \phi\left(a + j \frac{(b-a)}{s}\right) - \frac{s}{b-a} \int_a^b \phi(x) dx \right| \leq (M+1) \times \max_{a \leq x \leq b} \phi(x). \quad (21)$$

Observe now that y^n is a strictly increasing function of y in $[0, \infty)$ implying that g^n is extremized at exactly the same $\alpha \in [0, 1]$ as g . Since g is independent of n and analytic on the closed interval $[\epsilon, 1 - \epsilon]$, it follows that it has at most M extrema in $[\epsilon, 1 - \epsilon]$ for some constant M , and therefore so does g^n for all $n > 0$. Finally, since $g_{\max} > g(\alpha)$ for all $\alpha \neq \alpha_{\max}$ we get that for all sufficiently large n , $g_{\max}^n > n^{3/2} g_*^n$. Thus, using (21), we can rewrite (20) as

$$\begin{aligned} & \sum_{z=0}^n \binom{n}{z} F(z/n)^n \\ & < \frac{1}{\sqrt{2\pi n}} \frac{1}{\sqrt{\epsilon(1-\epsilon)}} \\ & \times \left(n \int_{\epsilon}^{1-\epsilon} g(\alpha)^n d\alpha + (M+2) g_{\max}^n \right). \end{aligned} \quad (22)$$

To deal with the integral in (22) we will use the Laplace method for asymptotic integrals. The following lemma can be found in [9, §4.2]:

Lemma 6 Let h be a real continuous function. Assume that there exist x_0 and $b, c > 0$ such that: i) $h(x) < h(x_0)$ if $x \neq x_0$, ii) $h(x) \leq h(x_0) - b$ if $|x - x_0| \geq c$, and iii) $h''(x_0) < 0$. If $\int_{-\infty}^{\infty} e^{h(x)} dx$ converges, then for any $\epsilon > 0$ and all sufficiently large t ,

$$\int_{-\infty}^{\infty} e^{th(x)} dx < \sqrt{\frac{2\pi}{(-h''(x_0) - 3\epsilon)t}} e^{th(x_0)}$$

and there is a similar lower bound for any $\epsilon < 0$.

To apply this lemma, we set $t = n$, and take any continuous h such that $h(x) = \ln g(x)$ for $x \in (0, 1)$, and such that $h(x)$ goes to $-\infty$ as $|x| \rightarrow \infty$ sufficiently fast so that $\int_{-\infty}^{\infty} e^{h(x)} dx$ converges. Observe that since $\ln y$ is strictly monotone in $[0, \infty)$, h is extremized at the same x as g . Clearly, condition ii) of Lemma 6 is also satisfied and since $[\ln g(x)]'' = g''(x)/g(x) - (g'(x)/g(x))^2$, we see that $h''(\alpha_{\max}) = g''(\alpha_{\max})/g(\alpha_{\max}) < 0$. Therefore, for all sufficiently large n

$$\begin{aligned} & \sum_{z=0}^n \binom{n}{z} F(z/n)^n \\ & < \frac{1}{\sqrt{2\pi n}} \frac{1}{\sqrt{\epsilon(1-\epsilon)}} \\ & \times \left(n \sqrt{\frac{2\pi g_{\max}}{-g''(\alpha_{\max})n}} g_{\max}^n + (M+2) g_{\max}^n \right) \\ & = C \times g_{\max}^n. \end{aligned}$$

In order to prove the lower bound, again we take $\epsilon = \min\{\alpha_{\max}, 1 - \alpha_{\max}\}/2 > 0$, and discard all the terms of the sum for which $\alpha \notin [\epsilon, 1 - \epsilon]$. Since $1/\sqrt{\alpha(1-\alpha)} \geq 2$, we have

$$\begin{aligned} \sum_{z=0}^n \binom{n}{z} F(z/n)^n & \geq \sum_{z=\epsilon n}^{(1-\epsilon)n} \binom{n}{z} F(z/n)^n \\ & > 2 \frac{36}{49} \frac{1}{\sqrt{2\pi n}} \sum_{z=\epsilon n}^{(1-\epsilon)n} g(z/n)^n. \end{aligned}$$

Replacing this sum by an integral as before and using the lower bound of Lemma 6 gives

$$\begin{aligned} & \sum_{z=0}^n \binom{n}{z} F(z/n)^n \\ & > \frac{72}{49} \frac{1}{\sqrt{2\pi n}} \\ & \times \left(n \sqrt{\frac{2\pi g_{\max}}{-g''(\alpha_{\max})n}} g_{\max}^n - (M+1) g_{\max}^n \right) \\ & = B \times g_{\max}^n. \end{aligned}$$